

CLAIMS

1. An apparatus for sharing input/output endpoints, the apparatus comprising:

sharing logic, coupled to a plurality of operating system domains through a load-store fabric, for routing transactions between said plurality of operating system domains and said sharing logic;
and

a first shared input/output endpoint, coupled to said sharing logic, configured to request/complete said transactions for said each of said plurality of operating system domains.
- 2 The apparatus as recited in claim 1, wherein said sharing logic routes said transactions according to a protocol that provides for routing of said transactions only for a single operating system domain, and wherein said sharing logic communicates with each of said plurality of operating system domains according to said protocol
3. The apparatus as recited in claim 1, wherein said load-store fabric comprises a point-to-point fabric.
4. The apparatus as recited in claim 1, wherein said sharing logic maps said each of said plurality of operating system domains to a corresponding root complex.

5. The apparatus as recited in claim 4, wherein said corresponding root complex is associated with one or more processing complexes, each comprising:

a processing core, configured to process data, wherein
said processing core accesses a memory that is
configured to store said data.
6. The apparatus as recited in claim 5, wherein said data is transferred from said corresponding root complex to said sharing logic according to a protocol that provides for routing of said transactions only for a single operating system domain.
7. The apparatus as recited in claim 5, wherein said data is transferred to said root complex from said sharing logic according to a protocol that provides for routing of said transactions only for a single operating system domain.
8. The apparatus as recited in claim 1, wherein said load-store fabric comprises hardware, software, or a combination of hardware and software that is configured to route said transactions.
9. The apparatus as recited in claim 1, wherein said transactions comprise memory read transactions, configured to transfer data to said plurality of operating system domains from said first shared I/O endpoint.

10. The apparatus as recited in claim 1, wherein said transactions comprise memory write transactions, configured to transfer data from said plurality of operating system domains to said first shared I/O endpoint.
11. The apparatus as recited in claim 1, wherein said transactions comprise message transactions, configured to direct said first shared I/O endpoint to perform a specified operation.
12. The apparatus as recited in claim 1, wherein said first shared input/output endpoint comprises a shared network interface controller (NIC).
13. The apparatus as recited in claim 12, wherein said shared NIC comprises a shared Ethernet controller.
14. The apparatus as recited in claim 1, wherein said first shared input/output endpoint comprises a Fiber Channel controller.
15. The apparatus as recited in claim 1, wherein said first shared input/output endpoint comprises a shared RAID controller.
16. The apparatus as recited in claim 1, wherein said first shared input/output endpoint is integrated into a shared I/O switch.
17. The apparatus as recited in claim 1, wherein said first shared input/output endpoint is integrated into a root complex.

18. The apparatus as recited in claim 1, wherein said first shared input/output endpoint is integrated into a processing complex.
19. The apparatus as recited in claim 1, wherein said protocol comprises PCI Express.
20. The apparatus as recited in claim 19, wherein said sharing logic isolates said transactions over a plurality of PCI Express bus hierarchies according to each of said plurality of operating system domains.
21. The apparatus as recited in claim 20, wherein a subset of said transactions are routed in accordance with addressing mechanisms for a particular one of said plurality of PCI Express bus hierarchies, said particular one of said plurality of PCI Express bus hierarchies corresponding to a particular one of said plurality of operating system domains.
22. The apparatus as recited in claim 1, wherein said sharing logic is configured to associate each of said transactions with a corresponding one of said plurality of operating system domains (OSDs), said corresponding one of said plurality of OSDs corresponding to one or more root complexes.
23. The apparatus as recited in claim 22, wherein said plurality of OSDs comprises 16 OSDs.

24. The apparatus as recited in claim 22, wherein said sharing logic designates said corresponding one of said plurality of OSDs according to a variant of a protocol that provides for routing of said transactions only for a single operating system domain.
25. The apparatus as recited in claim 24, wherein said variant comprises encapsulating an OS domain header within a transaction layer packet that otherwise comports with said protocol.
26. The apparatus as recited in claim 25, wherein said transaction layer packet is routed between said sharing logic and said first shared I/O endpoint.
27. The apparatus as recited in claim 26, wherein said first shared I/O endpoint is configured to detect said OS domain header and to perform a specified operation according to said protocol exclusively for said one of said plurality of OSDs.
28. The apparatus as recited in claim 1, wherein said shared I/O endpoint designates one of said transactions for a particular one of said plurality of operating system domains by encapsulating an OS domain header within a transaction layer packet, and wherein said transaction layer packet is routed to said sharing logic.

29. The apparatus as recited in claim 28, wherein said sharing logic is configured to detect and decapsulate said OS domain header from said transaction layer packet, and is configured to route said one of said transactions to said particular one of said plurality of operating system domains according to protocol that provides for routing of said transactions only for a single operating system domain.
30. The apparatus as recited in claim 1, wherein said sharing logic comprises a shared I/O switch.
31. The apparatus as recited in claim 1, wherein said sharing logic is integrated within a root complex.
32. The apparatus as recited in claim 1, wherein said sharing logic is integrated within a processing complex.
33. The apparatus as recited in claim 1, further comprising:
- a second shared input/output endpoint, coupled to said sharing logic, configured to request/complete additional transactions for a subset of said plurality of operating system domains.

34. The apparatus as recited in claim 33, wherein said sharing logic is configured to associate each of said additional transactions with one of a plurality of root complexes, said one of a plurality of root complexes corresponding to one of said subset of said plurality of operating system domains.

35. The apparatus as recited in claim 1, further comprising:

second sharing logic, coupled to said sharing logic and to a second shared input/output endpoint, configured route said transactions between said plurality of operating system domains and said second shared input/output endpoint.

36. The apparatus as recited in claim 1, further comprising:

second sharing logic, coupled to said sharing logic, a second shared input/output endpoint, and one of said plurality of operating system domains, configured route said transactions between said one of said plurality of operating system domains and said second shared I/O endpoint, and between said one of said plurality of operating system domains and said sharing logic.

37. The apparatus as recited in claim 1, further comprising:

second sharing logic, coupled to said sharing logic and one or more additional operating system domains, configured to route said transactions between said one or more additional operating system domains and said sharing logic.

38. A shared input/output (I/O) mechanism, comprising:

a load-store fabric, configured to enable operating system domains to share one or more I/O endpoints, said load-store fabric comprising:

sharing logic, coupled to said operating system domains and to said one or more I/O endpoints, for routing first transactions to/from said operating systems domains, and for routing second transactions to/from said one or more I/O endpoints, wherein each of said second transactions designates an associated one of said operating system domains for which an operation specified by each of said first transactions be performed.

39. The shared I/O mechanism as recited in claim 38, wherein , wherein said first transactions comport with a protocol that provides exclusively for a single operating system domain within said load-store fabric.

40. The shared I/O mechanism as recited in claim 38, wherein said load-store fabric comprises a point-to-point fabric.

41. The shared I/O mechanism as recited in claim 38, wherein said load-store fabric comprises a bus.
42. The shared I/O mechanism as recited in claim 38, wherein said first transactions comport with PCI Express.
43. The shared I/O mechanism as recited in claim 42, wherein said sharing logic isolates said first and second transactions over a plurality of PCI Express bus hierarchies according to each of said operating system domains.
44. The apparatus as recited in claim 43, wherein said first and second transactions are routed in accordance with addressing mechanisms for a particular one of said plurality of PCI Express bus hierarchies, said particular one of said plurality of PCI Express bus hierarchies corresponding to said associated one of said operating system domains.
45. The shared I/O mechanism as recited in claim 38, wherein said sharing logic associates each of said operating system domains with each of a corresponding root complex.
46. The shared I/O mechanism as recited in claim 45, wherein said corresponding root complex is associated with a processing complex that is configured to process data and to accesses a memory that is configured to store said data.

47. The shared I/O mechanism as recited in claim 46, wherein said data is transferred from said root complex to said sharing logic according to a protocol that provides exclusively for a single operating system domain within said load-store fabric.
48. The shared I/O mechanism as recited in claim 46, wherein said data is transferred to said root complex from said sharing logic according to a protocol that provides exclusively for a single operating system domain within said load-store fabric.
49. The shared I/O mechanism as recited in claim 38, wherein said load-store fabric comprises hardware, software, or a combination of hardware and software that routes said first and second transactions.
50. The shared I/O mechanism as recited in claim 38, wherein said operation comprises a memory read operation, directing that data be transferred from one of said one or more I/O endpoints to said associated one of said operating system domains.
51. The shared I/O mechanism as recited in claim 38, wherein said operation comprises a memory write operation, directing that data be transferred to one of said one or more I/O endpoints from said associated one of said operating system domains.
52. The shared I/O mechanism as recited in claim 38, wherein one of said one or more I/O endpoints comprises a shared network interface controller (NIC).

53. The shared I/O mechanism as recited in claim 52, wherein said shared NIC comprises a shared Ethernet controller.
54. The shared I/O mechanism as recited in claim 38, wherein said one of said one or more I/O endpoints comprises a Fiber Channel controller.
55. The shared I/O mechanism as recited in claim 38, wherein said one of said one or more I/O endpoints comprises a shared RAID controller.
56. The shared I/O mechanism as recited in claim 38, wherein one of said one or more I/O endpoints is integrated into a shared I/O switch.
57. The shared I/O mechanism as recited in claim 38, wherein one of said one or more shared I/O endpoints is integrated into a root complex.
58. The shared I/O mechanism as recited in claim 38, wherein one of said one or more shared I/O endpoints is integrated into a processing complex.
59. The shared I/O mechanism as recited in claim 38, wherein said sharing logic associates each of said first and second transactions with one of a plurality of root complexes, said one of a plurality of root complexes corresponding to one of said operating system domains.

60. The shared I/O mechanism as recited in claim 59, wherein said operating system domains comprise 16 operating system domains.
61. The shared I/O mechanism as recited in claim 38, wherein said second transactions comport with a variant of a protocol that provides exclusively for a single operating system domain within said load-store fabric.
62. The shared I/O mechanism as recited in claim 61, wherein said variant comprises encapsulating an OS domain header within a transaction layer packet of said each of said second transactions, wherein said each of said second transactions otherwise comports with said protocol.
63. The shared I/O mechanism as recited in claim 62, wherein said transaction layer packet is routed between said sharing logic and one of said one or more I/O endpoints.
64. The shared I/O mechanism as recited in claim 63, wherein said one of said one or more I/O endpoints is configured to detect said OS domain header and to perform said operation according to said protocol exclusively for said one of said operating system domains.
65. The apparatus as recited in claim 38, wherein said sharing logic comprises a shared I/O switch.

66. The apparatus as recited in claim 38, wherein said sharing logic is integrated within a root complex.
67. The apparatus as recited in claim 38, wherein said sharing logic is integrated within a processing complex.
68. The shared I/O mechanism as recited in claim 38, wherein one of said one or more I/O endpoints designates one of said second transactions for a particular one of said operating system domains by encapsulating an OS domain header within a transaction layer packet, and wherein said transaction layer packet is routed to said sharing logic.
69. The shared I/O mechanism as recited in claim 69, wherein said sharing logic is configured to detect and decapsulate said OS domain header from said transaction layer packet, and configured to route a corresponding one of said first transactions to said particular one of said plurality of operating system domains according to a protocol that provides exclusively for a single operating system domain within said load-store fabric.

70. The shared I/O mechanism as recited in claim 38, wherein one of said one or more I/O endpoints is configured to be shared with a first subset of said operating system domains, and wherein a second one of said one or more I/O endpoints is configured to be shared with a second subset of said operating system domains.
71. The shared I/O mechanism as recited in claim 38, further comprising:
- second sharing logic, coupled to said sharing logic and to second one or more I/O endpoints, configured route said first transactions, and configured to route third transactions to/from said second one or more I/O endpoints, wherein each of said third transactions designates a second associated one of said operating system domains for which a second operation specified by each of said first transactions be performed.
72. The shared I/O mechanism as recited in claim 38, further comprising:

second sharing logic, coupled to said sharing logic, second one or more I/O endpoints, and one of said operating system domains, configured route said first transactions, and configured to route third transactions to/from said second one or more I/O endpoints and to/from said sharing logic, wherein each of said third transactions designates a second associated one of said operating system domains for which a second operation specified by each of said first transactions be performed.

73. The shared I/O mechanism as recited in claim 38, further comprising:

second sharing logic, coupled to said sharing logic and one or more additional operating system domains, configured route third transactions between said one or more additional operating system domains and said sharing logic, said third transactions comporting with a protocol that provides exclusively for a single operating system domain within said load-store fabric, and configured to route fourth transactions to/from said sharing logic, wherein each of said fourth transactions designates a corresponding one of said one or more additional operating system domains for which a second operation is to be performed by said one or more I/O endpoints.

74. A method for sharing an input/output (I/O) endpoint within a load-store fabric, comprising:

via sharing logic within the load-store fabric,
interconnecting operating system domains to the
I/O endpoint;

first communicating with each of the operating system
domains individually; and

second communicating with the I/O endpoint according
to a variant of a protocol that provides
exclusively for a single operating system domain
within the load-store fabric to enable the I/O
endpoint to associate a prescribed operation with
a corresponding one of the operating system
domains.

75. The method as recited in claim 74, wherein said
interconnecting comprises:

coupling the operating system domains to the I/O
endpoint over a point-to-point fabric.

76. The method as recited in claim 74, wherein said
interconnecting comprises:

coupling root complexes to the I/O endpoint over a
bus.

77. The method as recited in claim 62, wherein said first
communicating comprises:

within the sharing logic, mapping each of the
operating system domains to a corresponding root
complex.

78. The method as recited in claim 77, wherein the corresponding root complex is associated with one or more processing complexes, each configured to process data, wherein the data is stored in a memory.
79. The method as recited in claim 78, further comprising:
transferring the data to from the corresponding root complex according to the protocol.
80. The method as recited in claim 78, further comprising:
transferring the data to the corresponding root complex according to the protocol.
81. The method as recited in claim 74, wherein said first communicating comprises:
issuing a memory read transaction directing that data be transferred to one of the operating system domains from the I/O endpoint.
82. The method as recited in claim 74, wherein said first communicating comprises:
issuing a memory write transaction directing that data be transferred from one of the operating system domains to the I/O endpoint.
83. The method as recited in claim 74, wherein said first communicating comprises:

issuing a message transaction directing that the I/O endpoint perform a prescribed operation associated with one of the operating system domains.

84. The method as recited in claim 74, wherein the I/O endpoint comprises a shared network interface controller (NIC).
85. The method as recited in claim 84, wherein the shared NIC comprises a shared Ethernet controller.
86. The method as recited in claim 74, wherein the I/O endpoint comprises a Fiber Channel controller.
87. The method as recited in claim 74, wherein the I/O endpoint comprises a shared RAID controller.
88. The method as recited in claim 74, wherein the I/O endpoint is integrated into a shared I/O switch.
89. The method as recited in claim 74, wherein the I/O endpoint is integrated into a root complex.
90. The method as recited in claim 74, wherein the I/O endpoint is integrated into a processing complex.
91. The method as recited in claim 74, wherein said first communicating comprises:

employing PCI Express as the protocol.
92. The method as recited in claim 91, further comprising:

isolating transactions over a plurality of PCI Express bus hierarchies according to each of the operating system domains.

93. The method as recited in claim 92, further comprising:

routing the transactions in accordance with addressing mechanisms for a particular one of the plurality of PCI Express bus hierarchies, the particular one of said plurality of PCI Express bus hierarchies corresponding to the corresponding one of the operating system domains.

94. The method as recited in claim 74, wherein said first communicating comprises:

associating the each of the operating system domains with each of a plurality of root complexes.

95. The method as recited in claim 94, wherein the plurality of operating system domains comprises 16 operating system domains.

96. The method as recited in claim 74, wherein said second communicating comprises:

employing the variant of the protocol to associate a unique root complex that with the corresponding one of the operating system domains.

97. The method as recited in claim 96, wherein said employing comprises:

encapsulating an OS domain header within a transaction layer packet that otherwise comports with the protocol, wherein the value of the OS domain header designates the corresponding one of the operating system domains.

98. The method as recited in claim 97, wherein said second communicating further comprises:

routing the transaction layer packet between the sharing logic and the I/O endpoint.

99. The method as recited in claim 98, wherein said second communicating further comprises:

within the I/O endpoint, detecting the OS domain header; and

performing the prescribed operation according to the protocol, wherein said performing is accomplished exclusively for the corresponding one of the operating system domains.

100. The apparatus as recited in claim 74, wherein the sharing logic comprises a shared I/O switch.

101. The apparatus as recited in claim 74, wherein the sharing logic is integrated within a root complex.

102. The apparatus as recited in claim 74, wherein the sharing logic is integrated within a processing complex.

103. The method as recited in claim 74, wherein said second communicating comprises:

within the I/O endpoint, designating a transaction for a particular one of the operating system domains by embedding an OS domain header within a transaction layer packet; and

first routing the transaction layer packet to the sharing logic.

104. The method as recited in claim 103, wherein said second communicating comprises:

within the sharing logic, detecting and removing the OS domain header from the transaction layer packet.

105. The method as recited in claim 104, wherein said first communicating comprises:

second routing the transaction layer packet to the particular one of the operating system domains according to the protocol.

106. The method as recited in claim 74, further comprising:

third communicating with a second I/O endpoint according to the variant of the protocol to enable the second I/O endpoint to associate specified operations with a corresponding one of a subset of the operating system domains.

107. The method as recited in claim 74, further comprising:

via second sharing logic within the load-store fabric,
second interconnecting additional operating
system domains to the sharing logic, whereby the
additional operating system domains can share the
I/O endpoint.

108. The method as recited in claim 74, further comprising:

via second sharing logic within the load-store fabric,
second interconnecting the operating system
domains to the sharing logic and to a second I/O
endpoint, whereby the operating system domains
can share both the I/O endpoint and the second
I/O endpoint.